

AD-A125 921

IMPLEMENTATION STRATEGY FOR AN INVENTORY FILTERING RULE
(U) NORTH CAROLINA UNIV AT CHAPEL HILL SCHOOL OF
BUSINESS ADMINISTRATION D BLAZER MAR 83 TR-25

1/1


UNCLASSIFIED

N00014-78-C-0467

F/G 12/1

NL

END
FILMED
-12-
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

ADA 125921


UNIVERSITY OF NORTH CAROLINA
AT CHAPEL HILL

IMPLEMENTATION STRATEGY FOR AN
INVENTORY FILTERING RULE

Technical Report #23

Douglas Blazer

March 1983



DTIC
ELECTE
MAR 22 1983
S E D

SCHOOL OF BUSINESS ADMINISTRATION
AND
CURRICULUM IN
OPERATIONS RESEARCH AND SYSTEMS ANALYSIS

This document has been approved
for public release and sale; its
distribution is unlimited.

DTIC FILE COPY

83 03 21 109

6

IMPLEMENTATION STRATEGY FOR AN
INVENTORY FILTERING RULE

Technical Report #25

Douglas Blazer

March 1983

Work Sponsored By

Office of Naval Research (N00014-78-C0467)

Decision Control Models in Operations Research

Harvey M. Wagner
Principal Investigator
School of Business Administration
University of North Carolina at Chapel Hill

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

DTIC
S 100

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report #25	2. GOVT ACCESSION NO. AD-A125921	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Implementation Strategy for An Inventory Filtering Rule		5. TYPE OF REPORT & PERIOD COVERED Technical
7. AUTHOR(s) Douglas J. Blazer		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of North Carolina at Chapel Hill Chapel Hill, N.C.		8. CONTRACT OR GRANT NUMBER(s) N00014-78-C-0467
11. CONTROLLING OFFICE NAME AND ADDRESS Mathematical and Information Sciences Division Office of Naval Research, Code 434 Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE March 1983
		13. NUMBER OF PAGES 20 and 7 (Appendix)
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) / In this report we examine the issues involved in implementing an inventory filtering rule that we showed performed well in Technical Reports #22 and #23. We discuss when and how to use the filtering rule and we examine the issue of non-stationary customer order distributions. We simulate a customer order distribution where the average ordinary customer order size has an increasing trend. We evaluate the statistical performance of the filtering rule for this non-stationary customer order distribution and find it continues to perform well.		

DD FORM 1473
1 JAN 73EDITION OF 1 NOV 66 IS OBSOLETE
S/N 0102-014-66011

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

FOREWARD

As part of the on-going program in "Decision Control Models in Operations Research," Mr. Douglas Blazer has concluded his study of removing large demands in the determination of an inventory replenishment policy. In this report, he examines the issues involved in the practical day-to-day use of an inventory filtering rule that identifies a threshold value T such that any order equal to or exceeding T is specially handled. The paper addresses the question: when and how to use the filtering rule, and what is the effect of non-stationary customer order distributions on the filtering rule. The paper provides the statistical results of using the filtering rule when the ordinary customer order size distribution is increasing. Other related reports dealing with this program are given on the following pages.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A	



- MacCormick, A. (1974), Statistical Problems in Inventory Control, ONR and ARO Technical Report 2, December 1974, School of Organization and Management, Yale University, 244 pp.
- Estey, A. S. and R. L. Kaufman (1975), Multi-Item Inventory System Policies Using Statistical Estimates: Negative Binomial Demands (Variance/Mean = 9), ONR and ARO Technical Report 3, September 1975, School of Organization and Management, Yale University, 85 pp.
- Ehrhardt, R. (1975), Variance Reduction Techniques for an Inventory Simulation, ONR and ARO Technical Report 4, September 1975, School of Organization and Management, Yale University, 24 pp.
- Kaufman, R. (1976), Computer Programs for (s,S) Policies Under Independent or Filtered Demands, ONR and ARO Technical Report 5, School of Organization and Management, Yale University, 65 pp.
- Kaufman, R. and J. Klinecicz (1976), Multi-Item Inventory System Policies Using Statistical Estimates: Sporadic Demands (Variance/Mean = 9), ONR and ARO Technical Report 6, School of Organization and Management, Yale University, 58 pp.
- Ehrhardt, R. (1976), The Power Approximation: Inventory Policies Based on Limited Demand Information, ONR and ARO Technical Report 7, June 1976, School of Organization and Management, Yale University, 106 pp.
- Klinecicz, J. G. (1976), Biased Variance Estimators for Statistical Inventory Policies, ONR and ARO Technical Report 8, August 1976, School of Organization and Management, Yale University, 24 pp.
- Klinecicz, J. G. (1976), Inventory Control Using Statistical Estimates: The Power Approximation and Sporadic Demands (Variance/Mean = 9), ONR and ARO Technical Report 9, November 1976, School of Organization and Management, Yale University, 52 pp.
- Klinecicz, J. R. (1976), The Power Approximation: Control of Multi-Item Inventory Systems with Constant Standard-Deviation-To-Mean Ratio for Demand, ONR and ARO Technical Report 10, November 1976, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 47 pp.
- Kaufman, R. L. (1977), (s,S) Inventory Policies in a Nonstationary Demand Environment, ONR and ARO Technical Report 11, April 1977, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 155 pp.

- Ehrhardt, R. (1977), Operating Characteristic Approximations for the Analysis of (s,S) Inventory Systems, ONR and ARO Technical Report 12, April 1977, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 109 pp.
- Schultz, C. R., R. Ehrhardt, and A. McCormick (1977), Forecasting Operating Characteristics of (s,S) Inventory Systems, ONR and ARO Technical Report 13, December 1977, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 47 pp.
- Schultz, C. R. (1979), (s,S) Inventory Policies for a Wholesale Warehouse Inventory System, ONR Technical Report 14, April 1979, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 75 pp.
- Schultz, C. R. (1980), Wholesale Warehouse Inventory Control with Statistical Demand Information, ONR Technical Report 15, December 1980, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 74 pp.
- Ehrhardt, R. and G. Kastner (1980), An Empirical Comparison of Two Approximately Optimal (s,S) Inventory Policies, Technical Report 16, December 1980, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 22 pp.
- Ehrhardt, R. (1980), (s,S) Policies for a Dynamic Inventory Model with Stochastic Lead Times, Technical Report 17, December 1980, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 20 pp.
- Mosier, C. (1981), Revised (s,S) Power Approximation, Technical Report 18, February 1981, School of Business Administration, University of North Carolina at Chapel Hill, 18 pp.
- Blazer, D. and M. McClelland (1981), An Inventory Model for Special Handling of Extreme Value Demands, Technical Report 19, December 1981, School of Business Administration, University of North Carolina at Chapel Hill, 10 pp.
- Mitchell, J.C. (1982), Choosing Single-Item Service Objectives in a Multi-Item Base-Stock Inventory System, Technical Report 20, School of Business Administration, University of North Carolina at Chapel Hill, 30 pp.

Blazer, D. (1983), Operating Characteristics for an Inventory Model That Special Handles Extreme Value Demand, Technical Report #21, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 15pp.

Blazer, D. (1983), Evaluation of a "Large Pop" Filtering Rule for Inventory Management Systems, Technical Report #22, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 32pp.

Blazer, D. (1983), Testing the Cost Effectiveness of an Inventory Filtering Rule Using Empirical Data, Technical Report #23, School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, 24pp.

ABSTRACT

In this report we examine the issues involved in implementing an inventory filtering rule that we showed performed well in Technical Reports #22 and 23. We discuss when and how to use the filtering rule and we examine the issue of non-stationary customer order distributions. We simulate a customer order distribution where the average ordinary customer order size has an increasing trend. We evaluate the statistical performance of the filtering rule for this non-stationary customer order distribution and find it continues to perform well.

TABLE OF CONTENTS

	<u>Page</u>
1. Introduction	1
2. When to Use the Inventory Filtering Rule	2
2.1 Alternative Replenishment Strategy	2
2.2 Filter for Spurious Demand	4
2.3 Allocation of Inventory Levels	5
3. General Data System Specifications	6
3.1 System Overview	6
3.2 Data Storage and Up-Dating Requirements	7
3.3 The Filtering Rule in a Non-Stationary Environment	11
3.3.1 Sensitivity to Non-Stationary Ordinary Order Distributions	13
3.3.2 Methods to Prevent Excessive Filtering	16
4. Areas for Future Research	18
References	20
Appendix	

1. INTRODUCTION

We have shown in [4], [3] that inventory systems that special handle extreme value demand can significantly reduce total expected cost. In [4], [3] extreme value demand was characterized in terms of the sum of all orders received during a week. In most practical situations, however, an inventory manager cannot suspend action on individual customer orders received during the week in order to determine if the total amount ordered for the week meets or exceeds some threshold value τ . A practical inventory filtering rule, therefore, must examine individual customer orders as they occur. In [2], [1] we describe an inventory filtering rule that identifies a point T such that any individual order equal to or exceeding T is specially handled. We hypothesized that using such an inventory filtering rule would result in truncated weekly demand distributions that, as was shown in [4], [3], significantly reduce total expected cost. Our hypothesis was supported in [2], [1], which demonstrated that the inventory filtering rule performed well from both a statistical and cost effectiveness standpoint.

In this report, we examine the issues involved in implementing an inventory filtering rule that special handles individual "large pop" customer orders. We address the questions: when and how to use the filtering rule, and what are the issues involved in implementing the filtering rule? We begin in Section 2 by describing some practical

settings for the use of the inventory filtering rule. In Section 3 we describe the general data system specifications for implementing the inventory filtering rule, including some data storage and updating requirements, and potential data processing problems. Finally, in Section 4, we discuss areas for further research.

2. WHEN TO USE THE INVENTORY FILTERING RULE

We have shown in [1] that the use of the inventory filtering rule can significantly reduce total expected cost. Here we address the issue of when the filtering rule should be used. And under what set of circumstances would the filtering rule provide significant cost savings? We describe three practical uses for the inventory filtering rule.

2.1 Alternative Replenishment Strategy

The first use of the rule is to fill the large (filtered) orders through an alternative replenishment strategy. Thus, ordinary orders are filled from existing stock and large orders are replenished via some special means. We cite three practical examples.

Example 1: Predict Large Orders

Assume large customer orders come from large customers. For example, consider a retail tire outlet that supports individual consumers but also supports three large customers, namely, two private trucking firms and a local vehicle rental agency. If these three large customers can predict their requirements or schedule their orders, thereby providing sufficient leadtime, the retail outlet can "special order" for its large customers. The cost savings resulting from stocking only for ordinary orders can be shared with the large customers to provide the incentive for these customers to participate in the scheme. Other factors, such as the price

and service provided by competitors, must of course also be considered. Provided that large customers can predict their requirements, the potential does exist for cost savings without a significant decrease in service.

Example 2: Use Other Procurement or Manufacturing Strategies

We cite a second example from the retail level of the United States Air Force inventory system. Normally, replenishment orders are sent to a central stocking point, which supports all retail outlets worldwide. For many items, leadtimes are relatively long. Note from [3], [1] that the potential for increased cost savings is greater the longer the lead-time. Thus, we propose that large customer orders be replenished locally at the retail level rather than ordered centrally. Local purchase of centrally stocked items is currently practiced in the United States Air Force inventory system for high priority requirements to reduce the lead-time. Although local purchase of items would probably incur an increased unit set-up cost, we note from [3], [1] that in many cases the breakeven special handling cost is 10 to 100 times the normal set-up cost. In this case, there is an added potential of reducing shipping and handling charges. Once again the potential exists for cost savings with only a small degradation in service. The degradation in service for special handling should be minimal due to the potential decreased leadtime for local purchase.

An analogous situation would hold for a manufacturing environment. For example, assume a make-to-order manufacturing company stocks component parts to reduce leadtime (and thereby stay competitive). This company could benefit by special handling of large orders. Special handling could mean increasing the leadtime for large orders or buying component parts

or end items from other sources. Depletion of existing component part stocks to fill a large order could put the company in a non-competitive position. Special handling of the large order would assure a smooth production process and a steady leadtime.

Example 3: Expedite Large Orders

Assume a multi-echelon inventory system where there are two levels: a central stocking point and small, numerous lower level outlets. Such a situation exists in businesses specializing in service and repair, such as computer or office machine companies. A problem inherent with such businesses is a large investment in inventory at the lower levels [6]. Passing large customer orders to the central stocking point rather than filling from existing local stock could significantly reduce inventory investment. The large orders could be handled expeditiously, for example, by a faster mode of transportation. The extra cost for transportation would be offset by the savings in inventory investment.

In all three examples, there is a potential for cost savings without a significant reduction in service. If the cost savings outweigh the alternative replenishment costs, some form of special handling should be used. We found in [3], [1] that for an inventory system with service levels of 80% or higher, and with a fixed set-up cost of 32 or less, that items with a holding cost per period equal to or greater than \$.50 showed the greatest cost savings and should definitely be considered for special handling.

2.2 Filter for Spurious Demand

The second use of the rule is to filter all orders to identify spurious large orders that are not likely to recur. We have seen [3],

[1] that include large pops in requirements computation significantly increases inventory investment. If these large orders are not likely to recur, then the resultant increased inventory investment is uneconomical. To illustrate, we have seen customer order data from an electronics company where large orders in one year were not repeated in the next year. Thus, in this setting the filtering rule would identify large orders and not include those orders in specifying future requirements. When filtering spurious demand in this manner, management has the option of using existing stock to fill the large order or to handle the order via some other means.

The filtering rule should be used for all items where an order is not likely to recur regardless of the item's value. The decision to include the large order in the determination of future requirements can be made via an on-line or off-line information processing mode. Thus, the filtering rule identifies the order as a large pop and the decision maker, presumably after consultation with the customer, must assess the likelihood of another order of that size recurring in the future.

2.3 Allocation of Inventory Levels

Assume that top management has concluded inventory levels are too high and that the inventory turnover rate is too low. Top management does not want to reduce service levels, yet inventory must be reduced. The filtering rule can reduce inventory investment without reducing the "nominal service level". If the nominal service level is 90%, then 90% of the "ordinary" demands are satisfied. Note from [3], [1] the significant amount of inventory investment that is saved versus the amount of demand left unfilled. For a relatively small reduction in service, there can be a significant decrease in inventory investment.

The filtering rule can be used to allocate scarce inventory investment funds as well. If inventory requirements total \$1000 and only \$750 is available for investment, the filtering rule can determine an objective way to allocate the \$750. Assuming some of the large pops will not recur, the filtering rule will perform better than an across the board cut or a reduction in the service level for all items.

3. GENERAL DATA SYSTEM SPECIFICATIONS

We now turn to how to use the rule. We discuss the general data specifications for the inventory filtering system. We begin in Section 3.1 by presenting an overview of the system. In Section 3.2 we describe the data storage and up-dating requirements, and in Section 3.3, we discuss a potential problem with the filtering rule.

3.1 System Overview

We seek to identify a point T such that any individual customer order of size T or greater is filtered. The value T is found using the filtering rule:

Let X_1, X_2, \dots, X_k be the k largest observed customer orders during the past N periods, where X_1 is the largest individual order and X_k the smallest. Given a value $r > 1$, let $X_0 = rX_1$ and define J_1 as the set j , for $1 \leq j \leq \rho k$ (for $\rho \leq 1$), such that $X_{j-1} \geq rX_j$. Given a value $\gamma > 0$, let $w = \gamma(X_1 - X_k)$. Define J_2 as the set of j , for $\rho k < j \leq k$, such that $X_{j-1} \geq rX_j$ and $X_{j-1} - X_j > w$. Set $T_r = \min_{j \in J_1 \cup J_2} (rX_j)$.

The parameters for the filtering rule are k , r , N , ρ , and γ . Parameter settings of $k=10$, $r=1.8$ and $\gamma=.2$ were tested and performed well [2], [1]. Set N such that there are between 25 to 50 orders in

N periods. We collect data to estimate the average number of orders per period and hence to determine N . If \bar{O}_N is the average number of orders in N periods, we set ρ by $\rho = .20 \frac{\bar{O}_N}{k}$. Thus, ρ is set so that the filtering rule includes an additional restriction ($X_{j-1} - X_j > w$) which must be met before filtering more than 20% of the customer orders.

For the sake of exposition, assume that the parameter values for k , r , and γ are as above, $N=26$ weeks, and $\bar{O}_{26}=25$; then $\rho=.5$. Further, assume that order history from the past 26 (N) weeks are available. Then to implement the filtering rule on a single item, the item's 10 (k) largest orders are selected and a value for T is found. The value T is used to filter out any large orders during the next 26 weeks. At the same time, collect the 10 largest orders (including filtered orders) for the next 26 weeks. Then reapply the filtering rule and find a new value for T . Thus for this example, a new value of T is found every 26 weeks.

We can simplify the filtering rule, especially if we can make some assumptions about the data. For instance, we can delete the parameter ρ and merely set a range depending on the parameter k . Thus, if $k=10$, we replace ρk with 5 and if $k=15$, we replace ρk with 9. Another simplification is to set $\rho=1$, which negates the need for γ . Note as \bar{O}_N approaches 50, ρ approaches 1 (for a given k). In essence, the more orders in a given period, the less there is a need for an additional restriction to filtering more than 20% of the orders.

3.2 Data Storage and Up-dating Requirements

Table 1 describes the data elements that must be stored to use the filtering rule.

The threshold value T requires on-line storage. Upon receipt of an order, the threshold value T is checked to determine if the order should be filtered. The value T is up-dated every N periods when the filtering rule program is run. The largest k customer orders X_1, \dots, X_k may be stored in an off-line file. At the end of the day, for each order received during the day, a check is made to determine if the current order is larger than any of the orders currently stored. If so, the current order size is stored and if there are k orders already in storage, the smallest order X_k is erased. Every N periods, after the filtering rule program is run, the values of X_1, \dots, X_k are reset to 0. If the values of X_1, \dots, X_k are not reset to 0, eventually the k largest observed customer orders will all be large pops and thus the rule will not filter any orders. Inventory managers may wish to save the previous interval's T value to compare to the new T value. If the T value changes by some percentage, say 100%, there may be a trend or shift in the order distribution that would require management action. We discuss non-stationary order distributions in the next section.

The Julian date of latest revision is the most recent date the filtering rule program was run and, hence, is the most recent date the value of T was changed. This date is checked every order to determine if N periods have passed since the latest revision (see Figure 1). If N periods have passed, the filtering rule program is called. Note it may be necessary for a second check to be made prior to calling the filtering rule program. The filtering rule can behave poorly if there are too few orders (for example 15 orders). Hence, if N periods have passed since T was previously revised, a check should be made of O_n .

the number of orders in n periods since the latest revision of T . If O_n is greater than some value, say 21, then the filtering rule program is called. If O_n is not greater than 21, then the filtering rule is not called until receipt of order number 22. The value O_n then should be stored on-line and incremented upon receipt of each order. When the filtering program is called, it resets O_n to zero.

DATA STORAGE REQUIREMENTS

Data Element	Description	Frequency of Access	Information Processing Model
T	Order threshold value	1) Check every order 2) Update every N periods	On-line
X_1 through X_k	Largest k customer order in N periods	1) Check and revise if necessary for every order via end-of-day batch processing 2) Restart every N periods	Off-line
Julian date of last review	Last date the value of T was updated	1) Check every order to determine if revision is necessary 2) Update every N periods	On-line
O_n	Number of orders since the last revision of T	1) Revise every order 2) Restart every N periods	On-line
N	Number of periods in revision interval	1) Check and revise if necessary every N periods	On-line
k, p, r, y	Parameters for the filtering rule	1) Revise via off-line management decision	Off-line in the filtering rule program

Table 1

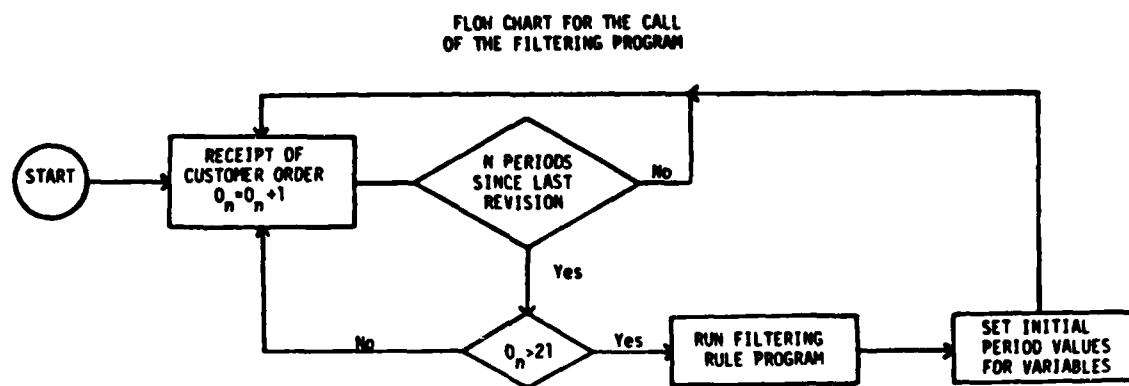


Figure 1

The filtering program, by means of a look-up table, can also be used to reset N if necessary. The first step to reset N is to determine \bar{O}_N , the average number of orders in N periods:

$$\bar{O}_N = \frac{O_N}{N} .$$

Then a look-up table determines N by:

- (1) If $.48 \leq \bar{O}_N < .95$, then $N=52$,
- (2) If $.95 \leq \bar{O}_N < 2.0$, then $N=26$,
- (3) If $2.0 \leq \bar{O}_N < 4.0$, then $N=13$,
- (4) If $4.0 \leq \bar{O}_N$, then $N=10$.

If there are fewer than 25 orders a year ($\bar{O}_N < .48$), N greater than 52 can be used. But caution should be exercised for longer revision intervals, especially if demand is subject to trends or shifts. We discuss the problem of demand shifts and trends in the next section.

The data elements k , ρ , r and γ are the parameters for the filtering rule. They are determined off-line by management and used in the filtering rule program. In Table 2, we provide a decision logic table for

changing the parameters of the filtering rule. For example, if management feels too many orders are being specially handled, then k or ρ can be decreased and r , γ and N can be increased or any combination of these actions. However, we found changing r and N to be the most effective parameters to alter [3].

DECISION LOGIC TABLE FOR CHANGING
THE FILTERING RULE PARAMETERS

If one increases the parameter,	then the frequency of filtering orders is:
k	increased
ρ	increased
r	decreased
γ	decreased
N	decreased

Table 2

3.3 The Filtering Rule in a Non-Stationary Environment

For the research in [4], [3], [2], [1] we have assumed a stationary distribution for customer orders. Suppose customer orders are non-stationary; for example, suppose the average customer order size is increasing. For an increasing trend in the average customer order size, there is a danger of filtering too many orders. If the customer order size is stationary, an increase in the number of demands per period increases demand per period but does not pose any problem for the filtering rule. In many cases, filtering too many orders tends to decrease the amount of cost savings [3] and may also make the enterprise less

competitive.

The danger of filtering too few demands is not so serious a problem as filtering too many demands. Note from [1] in a number of instances fewer than 5% of customer orders were filtered and significant cost savings were still achieved. Thus, we seek to limit the filtering of too many orders.

To examine the issue of non-stationary order distributions, we categorize orders into ordinary orders (those orders that are usually not filtered) and large orders (those that are normally filtered). There are three possible trends for each type of orders: increasing, stable, and decreasing. The nine possible combinations are shown in Table 3.

COMBINATIONS FOR TRENDS IN THE
AVERAGE ORDER SIZE

If the ordinary order trend is:	and the large order trend is:	then the overall trend is:	and the probability of excessive filtering is:
Increasing	Increasing Stable Decreasing	Increasing Increasing Indeterminate	Indeterminate Increasing Increasing
Stable	Increasing Stable Decreasing	Increasing Stable Decreasing	Decreasing Stable Decreasing
Decreasing	Increasing Stable Decreasing	Indeterminate Decreasing Decreasing	Decreasing Decreasing Decreasing

Table 3

Using the potential for filtering too many orders as the criterion, Table 3 shows that a problem is likely to occur only if the average size of an ordinary order is increasing. Thus, if there is a potential for a significant increase in the size of ordinary orders, inventory managers should take some precautions in using the filtering rule.

3.3.1 Sensitivity to Non-stationary Ordinary Order Distributions

We next examine how sensitive the filtering rule is for non-stationary ordinary order distributions. We use a simulated customer order distribution from [2] with proportion P of the order distribution as ordinary orders, and the remaining $1-P$ as large pops. We introduce an increasing two-year trend to the ordinary order distribution. The basic order distribution is described below. Let ϕ_z be the probability of a customer order size z , where

$$\phi_z = \begin{cases} \phi(1.5) & z=1 \\ \phi(z+.5) - \phi(z-.5) & z=2, \dots, B-1 \\ \phi(b) - \phi(z-.5) & z=B \\ \frac{1-P}{I} & z \in Z \end{cases}$$

where

$$\phi(z) = \int_0^z \lambda_n e^{-\lambda_n z} dz,$$

$$b = \frac{\ln(1-P)}{-\lambda_0},$$

$$B = \text{integer}(b+.5), \quad (1)$$

$$B+ = \left\lceil \frac{B+.5}{10} \right\rceil * 10,$$

$$Z = \{z | z=B+(10)100(100)1000\},$$

$$I = \text{dimension of } Z.$$

We let $P=.90$, $\lambda_0=.1$, $b=29.957$, $B=30$, $B+=40$, $Z=40, 50, \dots, 100, 200, \dots, 1000$ and $I=16$. We set λ_n to be

$$\frac{1}{\lambda_n} = \frac{1}{\lambda_0} + \frac{n\ell}{\lambda_0 104} \quad \text{for } n=0 \text{ to } 104, \quad (2)$$

where λ is the proportion increase for two years combined. We set $\lambda=.40$, 1.60, and 2.00. Thus, $\lambda=.40$ means there is a 40% increase in the small order size distribution in two years or 20% a year, $\lambda=1.60$ means a 160% increase, and $\lambda=2.00$ means a 200% increase.

We generate orders to apply to the filtering rule. The six sets of parameters for the filtering rule are shown in Table 4.

FILTERING RULE PARAMETERS						
Case No.	N	k	γ	r	ρ	λ
1	26	10	.2	1.8	.5	.40
2	26	10	.2	1.8	.5	1.60
3	26	10	.2	1.8	.5	2.00
4	56	15	.2	1.8	.67	.40
5	52	15	.2	1.8	.67	1.60
6	52	15	.2	1.8	.67	2.00

Table 4

We use two levels for N. We generate .96 orders per week, so on the average the filtering rule uses 25 orders for N=26 and 50 orders for N=52.

We generate N periods of random customer orders and apply the filtering rule to find a value for T. We then generate another N periods of customer orders and find the associated value for T. For all cases we generate 5000 weeks of orders. We reset the value of λ_n to λ_0 every two years (104 weeks). Thus, for N=26, we generate $48 \left(\frac{5000}{104} \right)$ two-year intervals and $192 \left(\frac{5000}{26} \right)$ values of T. For N=52, we generate 48 two-year intervals and $96 \left(\frac{5000}{52} \right)$ values of T.

In Appendix I we show the distribution of the T values for the six cases. We summarize the results in Table 5 and compare them to the

results we found in [2] for the stationary customer order size distribution with the same parameters. Note as N increases, less customer orders are filtered, the same results as we saw for the stationary customer order distribution. Even for a yearly 100% increase ($\lambda=2.00$) in the average customer order size, there is only a small probability of filtering more than 20% of the orders in the $N=26$ case and no possibility for the $N=52$ cases. Note that the T values increase with the trend and the rule actually filters a smaller average percentage of orders than it did for the stationary distribution; however, the variance of the T values is larger.

SUMMARY OF RESULTS FOR CASES 1 THROUGH 6

Case Number	N	λ	Average Percent of Orders Filtered	Variance of Percent of Orders Filtered	Relative Frequency of Filtering $>.20$
1	26	.40	8.55	1.76	.0
2	26	1.60	9.13	9.66	.005
3	26	2.00	7.49	12.39	.010
4	52	.40	7.51	2.43	.0
5	52	1.60	6.24	1.98	.0
6	52	2.00	5.63	1.76	.0
Stationary Customer Order Distribution	25		8.95	3.30	.005
	50		7.93	1.29	.0

Table 5

We examine two additional cases, where there is no increase in the customer order size for the first N periods and then an 80% increase ($\lambda=1.60$) per year for the remaining periods in the 2-year interval. Thus we change (2) to

$$\frac{1}{\lambda_n} = \begin{cases} \frac{1}{\lambda_0} & \text{for } n = 0 \text{ to } N \\ \frac{1}{\lambda_0} + (n-N) \frac{\ell}{\lambda_0 104} & \text{for } n = N+1 \text{ to } 104. \end{cases} \quad (3)$$

For each case 7 we use $N=26$, $k=10$, $\rho=.5$ and for case 8, we use $N=52$, $k=15$, $\rho=.67$. We use $\gamma=.2$ and $r=1.8$ for both cases.

We display the results for cases 7 and 8 in Appendix I and we summarize the findings in Table 6. Again we compare the results to the results with the stationary customer order size distribution we found in [2]. Even for cases where small order size begins increasing immediately after updating T , there is only a small probability of filtering more than 20% of the orders for $N=26$ and no chance with $N=52$. Again, the results are comparable to the results for the stationary customer order distribution; only the variance of T values are larger.

SUMMARY OF RESULTS FOR CASES 7 AND 8
(NO TREND IN INITIAL N PERIODS)

Case Number	N	λ	Average Percent of Orders Filtered	Variance of Percent of Orders Filtered	Relative Frequency of Filtering $>.20$
7	26	1.60	9.86	11.63	.010
8	52	1.60	7.79	3.30	.0
Stationary Customer Order Distribution	25		8.95	3.30	.005
	50		7.93	1.29	.0

Table 6

3.3.2 Methods to Prevent Excessive Filtering

Thus, non-stationary demand is not a major obstacle to applying the filtering rule. If small order sizes are increasing, the inventory manager should increase N , thereby increasing the number of orders used

to find T . It may not always be practical to increase N . For example, if the average number of orders per year is 25, an inventory manager may not wish to wait 2 years prior to updating T . For these situations, we recommend three methods to prevent excessive filtering.

The first method is to use a rolling horizon. Two years of orders, call them year 1 and year 2, might be used to find a value for T . At the end of year 3 a new value for T would be found using data from year 2 and year 3. Thus, T is updated yearly using the most current two years of data. This method requires storage of the dates of the orders so as to determine when the data are outdated and should be replaced.

Another method to prevent filtering too many orders is to set an upper value for the number of orders filtered in N periods and if that value is reached, off-line management action is taken. This requires storage of a new data element, call it F_n , which is the number of orders that have been filtered in n periods (for $n=1$ to N). Thus, whenever F_n equals some percentage of the average number of orders in N periods, the item is flagged for off-line management action. For example, suppose management wants to ensure that no more than 20% of the orders are filtered. If the average number of orders in N periods (\bar{O}_N) is 50, then whenever $F_n=10$ management is notified. Since there may not be enough order data collected to call the filtering program, management would determine the new T value based on available data. The problem with this procedure is that there still may be too many orders filtered. Also, it should not be necessary to wait for all 10 orders to be filtered before taking corrective action.

A third method to prevent excessive filtering employs binomial sampling techniques [5]. For this method, count the number of filtered orders in n periods (for $n=1$ to N). Assume that the number of orders occurring during any n interval is O_n , and that at most, some proportion p of the orders, say 20%, are considered large. Then p and O_n are parameters for a binomial distribution. We can test the hypothesis

$$H_0: p \leq .20 \text{ versus } H_1: p > .20$$

for some value of O_n and for some significance level. We can find an acceptance number that denotes the maximum number of large orders that can be found and still accept the null hypothesis. Look-up tables can be prepared for O_n values of 2 through 50. For example, if $O_n=10$, then the acceptance number equals 4 at the significance level of 10%. Thus, if the tenth order is the fifth order to be filtered, then reject the null hypothesis, and conclude the rule is filtering more than 20% of the orders. The system then may take two actions; it can flag the item for off-line management action or if enough customer orders have been collected it can revise the value T .

4. AREAS FOR FUTURE RESEARCH

We suggest three areas for future research. First, the problem of filtering orders from a non-stationary order distribution needs further research. In the discussion in Section 3.3, we assumed that the demand shift is uncertain, and we sought only to prevent excessive filtering of orders. If the shift is predictable, however, the filtering rule could be modified. For example, the rule could be modified to accomodate

seasonality.

A second area for research is to examine a periodic review, infinite horizon inventory model with linear holding and penalty cost and a fixed set-up cost. We seek to determine an optimal threshold value τ such that any period's demand equal to or exceeding τ is specially handled. We examined in [4] an optimal threshold value for the newsboy problem. An algorithm to provide an optimal τ value for (s,S) policies could lead to additional insights, which could aid users of the filtering rule or help in the development of new filtering rules.

A third area for further research is to examine the case of lost demand (sales). In this research, we assumed that all unfiltered orders were met or backordered and all filtered orders were specially handled. If we assume instead that all unfiltered are not issued from stock and are lost as are all filtered orders, what is the impact of filtering?

REFERENCES

1. Blazer, Douglas J., Technical Report #23 "Testing the Cost Effectiveness of an Inventory Filtering Rule Using Empirical Data," School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, February, 1983.
2. Blazer, Douglas J., Technical Report #22, "Evaluation of a 'Large Pop' Filtering Rule for Inventory Management Systems," School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, February, 1983.
3. Blazer, Douglas J., Technical Report #21, "Operating Characteristics for an Inventory Model that Special Handles Extreme Value Demands," School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, January, 1983.
4. Blazer, Douglas J. and M. McClelland, Technical Report #19, "An Inventory Model for Special Handling Extreme Value Demands," School of Business Administration and Curriculum in Operations Research and Systems Analysis, University of North Carolina at Chapel Hill, December, 1981.
5. Pfaffenberger, Roger C., and J. H. Patterson, Statistical Methods for Business and Economics, Homewood, Illinois: Richard D. Irwin, Inc., 1977.
6. Smith, Stephen A., Executive at the Xerox Research Center, Palo Alto, California, Interviews and Presentation conducted at the Multi-Echelon Inventory Systems Conference at the University of Santa Clara on June 14 and 15, 1982.

APPENDIX

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-22	.0	-
23	.010	.141
26	.005	.114
28	.005	.104
30	.005	.094
32	.021	.093
34-39	.230	.092
41-49	.377	.087
50-59	.188	.082
71	.068	.068
89	.026	.061
107-180	.063	.048
≥200	.0	-

SIMULATION RESULTS FOR CASE 1

N=26 $\lambda=.40$

Table A-1

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-27	.0	-
28	.005	.204
32	.016	.167
34	.005	.152
36	.010	.140
37	.010	.135
39	.021	.126
41	.005	.114
43	.037	.108
44	.037	.104
46	.026	.100
48	.026	.096
50	.031	.089
52	.021	.087
54	.042	.086
55	.021	.085
57	.016	.084
59	.031	.084
61	.037	.075
62	.016	.075
64-69	.052	.075
70-79	.188	.068
80-89	.047	.061
90-94	.063	.055
100-180	.236	.048
≥200	.0	-

SIMULATION RESULTS FOR CASE 2

N=26 $\lambda=1.60$

Table A-2

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-25	.0	-
26	.005	.259
28	.005	.236
35	.021	.173
37	.016	.162
39	.010	.148
41	.010	.133
43	.021	.125
44	.021	.121
46	.031	.116
48	.016	.111
50	.026	.100
52	.021	.096
54	.047	.092
55	.021	.090
57	.010	.087
59	.016	.086
61	.016	.076
62	.021	.076
64	.026	.075
66	.016	.074
68	.016	.074
70-79	.131	.066
80-89	.089	.059
90-99	.079	.653
100-180	.309	.047
≥200	.0	-

SIMULATION RESULTS FOR CASE 3

N=26 $\lambda=2.00$

Table A-3

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-33	.0	-
34-39	.085	.092
41-49	.347	.087
50-56	.242	.083
71-72	.084	.068
89	.032	.061
107-180	.211	.049
≥200	.0	-

SIMULATION RESULTS FOR CASE 4

N=52 $\rho=.40$

Table A-4

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-43	.0	-
44	.011	.100
50	.011	.087
52	.011	.085
54	.032	.084
55	.042	.084
57-59	.084	.083
61-69	.094	.075
70-79	.147	.068
80-89	.105	.061
90-99	.021	.055
100-180	.442	.049
≥200	.0	-

SIMULATION RESULTS FOR CASE 5

N=52 $\rho=1.60$

Table A-5

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-49	.0	-
50	.011	.096
54	.042	.091
55	.011	.090
59	.011	.086
61	.032	.077
62	.011	.076
64	.032	.075
66	.021	.075
71-79	.116	.068
80-89	.084	.061
90-100	.042	.055
101-160	.589	.049
≥200	.0	-

SIMULATION RESULTS FOR CASE 6

N=52 $\lambda=2.00$

Table A-6

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-16	.0	-
17	.005	.300
23	.005	.206
25	.005	.183
26	.010	.183
28	.016	.152
30	.010	.133
32	.031	.123
34	.021	.117
36	.050	.103
37	.021	.106
39	.084	.102
41	.042	.093
43	.037	.090
44	.037	.089
46	.010	.088
48	.005	.087
50	.010	.082
52	.016	.081
54-59	.141	.061
61-69	.105	.073
70-76	.131	.067
84-89	.053	.060
91-94	.010	.054
107-180	.131	.048
≥200	.0	-

SIMULATION RESULTS FOR CASE 7

N=26 $\lambda=1.60$

Table A-7

T Value	Relative Frequency of T Value	Cumulative Relative Frequency of Filtered Orders
1-31	.0	-
32	.021	.101
34	.021	.097
35	.032	.096
37	.042	.094
39	.126	.093
40-49	.126	.087
50-59	.168	.083
61-67	.032	.075
70-72	.112	.068
89	.053	.061
107-180	.263	.049
≥200	.0	-

SIMULATION RESULTS FOR CASE 8

N=52 $\lambda=1.60$

Table A-8

